Sri Lanka's Social Media Ban Won't Stop Predators. Verified Identity and Prosecution Will.

25th of August 2026

10.13140/RG.2.2.35429.15845

ResearchGate

Anthony Vondoom
ORCID: 0009-0003-4953-1427
www.dssm.academy
anthony.vondoom@outlook.com

Abstract

Sri Lankan children are being groomed, sexually extorted, and exploited online right now, and Parliament's answer — a minimum age of 16 for social media — checks a birthdate, not a person. This article shows why that fails using the country's own evidence: National Child Protection Authority complaint data, global threat assessments, an original investigation the author reported to Sri Lankan authorities and to the U.S. National Center for Missing & Exploited Children (NCMEC), and Sri Lanka's own legal-gap analysis showing grooming isn't even a codified crime yet. Australia already ran this experiment; its own government evaluation found more than 85% of under-16s still on banned platforms three months in, because a self-declared birthdate stops nobody. The harm behaves as a loop — unverified contact causes injury, injury increases vulnerability to the next unverified contact — and a ban that never verifies who is actually messaging a child does not touch either end of it. Closing the loop takes two things a minimum-age law does neither: prosecuting the adults responsible under a strengthened legal code, and replacing anonymous accounts with verified-identity accountability applied to everyone, not just children. Sri Lanka is not unique. It is simply where this fight is happening loudly enough, right now, to show every other country what works and what doesn't.

Keywords

Sri Lanka social media ban; online child sexual exploitation and abuse; child grooming prevention; social media age verification; verified identity accountability; child protection policy; online predators; technology-facilitated child abuse

Sustainable Development Goals

This article's findings and recommendations are relevant to the following UN Sustainable Development Goals.

SDG 16 (Peace, Justice and Strong Institutions), and specifically Target 16.2 — end abuse, exploitation, trafficking, and all forms of violence against and torture of children — is addressed directly by this article's central argument for offender prosecution and for closing the legislative gap on grooming and computer-generated CSAM documented in Section 3.3.

SDG 5 (Gender Equality), and specifically Target 5.2 — eliminate all forms of violence against women and girls in public and private spheres — is addressed through the article's engagement with the technology-facilitated gender-based violence data reviewed in Section 3.1.

SDG 3 (Good Health and Well-being) is addressed through the article's engagement with the neurodevelopmental and trauma-psychological evidence reviewed in Section 3.5 and Section 4.2.

SDG 17 (Partnerships for the Goals) is reflected in the article's reliance on, and advocacy for strengthening, the cross-border reporting and enforcement pathways — INTERPOL, NCMEC, and the Internet Watch Foundation — documented in Section 3.2.

1. Introduction

A child in Sri Lanka does not have to leave their bedroom to be groomed, sexually extorted, or exploited for images that will circulate for the rest of their life. Parliament's answer, so far, is a bill that checks a birthdate. On 8 July 2026, Prime Minister Harini Amarasuriya told Parliament that the Ministry of Women and Child Affairs had opened preliminary stakeholder talks on restricting or banning social media access for children under 16 — a jump from the under-12 threshold the Ministry had floated as recently as January 2026. A private member's bill from MP Faizer Mustapha is moving through procedural stages that would give the country its first dedicated legal minimum age for social media use, placing it alongside Australia, France, and the United Kingdom. A draft government bill, the Online Protection for Children Bill, was reported heading to Cabinet in late July 2026, though no public text of it exists to date.

The harm behind this legislative activity is usually discussed in the language of risk, exposure, and correlation. That language understates it. Global and domestic evidence, reviewed in Section 3, documents active online child sexual exploitation and abuse (OCSEA) affecting Sri Lankan minors right now: grooming, sexual extortion (sextortion), and the production, possession, and circulation of child sexual abuse material (CSAM), alongside coordinated technology-facilitated violence against

women and girls. These are not adjacent concerns to a neurodevelopmental argument about screen time. They are criminal acts, presently occurring, against a population unable to consent to the risk or to defend itself against it. The adults committing these acts must be identified, prosecuted, and dealt with under the law — a criminal-justice obligation that no platform-access policy satisfies on its own. What platform policy can and should do is close the one structural gap that makes all of it possible: an unidentified adult can currently reach a child directly, repeatedly, and at scale, without ever having to say who they are.

This article addresses two linked research questions. First, does the available evidence support a blanket, identity-verified minimum-age ban as the appropriate policy instrument for the harm documented against Sri Lankan minors? Second, given that evidence, what mechanism would actually reduce the harm rather than merely relocate or symbolically address it? Existing literature establishes the scale of the underlying harm without resolving the second question. Institutional and survey data (NCPA, 2026; UNICEF, 2018; Save the Children, cited in EconomyNext, 2024; UN Women & UNFPA, 2025) and global OCSEA threat assessments (WeProtect Global Alliance, 2025; ECPAT International, INTERPOL, & UNICEF Innocenti, 2022) establish that online harm to Sri Lankan children and women, including criminal sexual exploitation, is extensive, active, and ongoing. Sri Lanka's own legal-gap analysis (Verité Research, 2019) documents that grooming is not codified as a distinct offence and that the Penal Code does not address computer-generated CSAM, leaving prosecutable conduct without a clean statutory basis. Neurodevelopmental literature (Paulus et al., 2019; Hutton et al., 2020; Maza et al., 2023) and trauma psychology (Finkelhor & Browne, 1985) establish mechanisms by which this exposure compounds over time rather than resolving after a single incident. Domestic policy analysis (Hattotuwa, 2026a, 2026b, 2026c) evaluates the specific ban mechanism against comparative and infrastructural evidence and finds it unsupported. No existing publication combines all of these strands into a single assessment of what Sri Lanka's proposed legislation should actually do. This article does so.

The thesis advanced is as follows. The underlying harm is real, active, criminal in its most serious forms, and urgent. It behaves as a recursive loop: unverified exposure produces psychological injury, which measurably increases vulnerability to further victimization, which in turn enables further unverified exposure. Breaking this loop requires two concurrent interventions operating on different points of the cycle — offender identification and prosecution under a strengthened legal framework, which removes the perpetrator; and mandatory verified-identity account accountability, applied to all users as a condition of holding an account rather than as an age-exclusion mechanism aimed only at minors, which removes the anonymity that allows unverified contact to occur in the first place. A blanket age ban enforced by self-declared birthdate, the mechanism currently before the Sri Lankan Parliament,

addresses neither point effectively, and the comparative evidence in Section 3 shows precisely why. The remainder of the article proceeds as follows. Section 2 describes the documentary and investigative methodology. Section 3 presents findings from institutional, survey, legal, investigative, and neurodevelopmental sources without interpretation, including two summary tables. Section 4 discusses these findings, models the recursive harm loop, situates the findings against the domestic policy debate, states limitations, and sets out the two-track recommendation in detail.

2. Methodology

This study uses a mixed-methods documentary and policy-analysis design combining six data strands, selected for their capacity to triangulate a question that no single dataset can answer alone.

First, secondary analysis of published institutional statistics: NCPA child abuse complaint data for January–June 2026 and the 2010–2025 decadal total, as reported to the Sri Lankan Parliament; the UNICEF (2018) national baseline study on child internet use; Save the Children survey data as reported by EconomyNext (2024); and the UN Women/UNFPA (2025) study on technology-facilitated gender-based violence in Sri Lanka. These sources were accessed in their published form; no primary statistical analysis was performed on the underlying microdata, which was not available to the author.

Second, analysis of global online child sexual exploitation and abuse (OCSEA) threat assessments and cross-country research: the WeProtect Global Alliance (2025) Global Threat Assessment, produced with Columbia University's CPC Learning Network, and the Disrupting Harm research programme conducted by ECPAT International, INTERPOL, and the UNICEF Office of Research – Innocenti (2022) across thirteen countries. These sources were used for global and regional benchmarking; no country-specific Sri Lankan OCSEA prevalence dataset of comparable scale currently exists, a gap noted explicitly in Section 4.3.

Third, analysis of Sri Lanka's domestic legal and regulatory framework governing online child sexual exploitation, drawing on the Computer Crimes Act No. 24 of 2007, Sri Lanka's 2005 accession to the Council of Europe Convention on Cybercrime (Budapest Convention), the Online Safety Act No. 9 of 2024, and an independent legal-gap analysis commissioned by ECPAT Sri Lanka (Verité Research, 2019).

Fourth, a targeted review of peer-reviewed neurodevelopmental and trauma-psychology literature: studies using the Adolescent Brain Cognitive Development (ABCD) Study cohort (Paulus et al., 2019; Hutton et al., 2020; Maza et al., 2023), a Sri Lanka-specific clinical study of internet addiction disorder (Ratnayake & Hewapathirana, 2023), a Sri Lankan psychiatric perspective piece (Sundararaj et al., 2025), and the foundational traumagenic dynamics model of child sexual abuse

(Finkelhor & Browne, 1985), used in Section 4.2 to ground the recursive-loop model in established clinical theory rather than metaphor.

Fifth, analysis of the Sri Lankan legislative and policy record on the proposed ban: parliamentary statements reported in Hansard and national media between January and August 2026, the text and status of the private member's bill on social media age restriction, and the reported status of the draft Online Protection for Children Bill.

Sixth, an original open-source investigation (OSINT) conducted by the author into the operating patterns of accounts and networks exploiting minors and women online in Sri Lanka, with findings submitted directly to Sri Lankan authorities, including the National Child Protection Authority and the police, and internationally to the U.S. National Center for Missing & Exploited Children (NCMEC), in 2026; and analysis of an independent third-party investigative report (Hattotuwa, 2026a) documenting platform-accountability failures on Meta's services in Sri Lanka, together with a formal policy brief and two published critiques by the same author addressing the proposed ban mechanism directly (Hattotuwa, 2026b, 2026c). Consistent with standard practice for investigative child-safety work, this article withholds the specific methods, platforms, and identifying details of the author's own investigation, since publication of such detail would assist the actors it describes rather than efforts to constrain them; only the structural finding relevant to this analysis — the absence of functioning identity or age verification on the networks observed — is reported. This is acknowledged as a methodological limitation in Section 4.3.

Across all six strands, the analytical approach is comparative and triangulating rather than statistical: findings from independent sources are cross-referenced for convergence or contradiction, and no new statistical model, survey, or experiment was conducted by the author. This is a documentary and policy-analytic design, not an empirical field study or a legal filing, and its results should be read accordingly. The article follows the conventional Introduction–Methods–Results–Discussion (IMRaD) structure by convention, for ease of navigation against the empirical literature it cites; Section 3, headed ‘Results’ in keeping with that convention, reports a narrative evidence synthesis rather than results in the sense of primary data collected by the author, and should be read as such.

. Results

Findings are reported by source strand, without interpretation; interpretation is reserved for Section 4. Terminology follows the usage of the cited sources: OCSEA

(online child sexual exploitation and abuse) is used as the umbrella term, with CSAM (child sexual abuse material) and grooming referring to specific, legally defined conduct within it.

3.1 Institutional complaint and survey data

The NCPA recorded 5,523 child abuse complaints in the first six months of 2026, comprising 213 sexual harassment complaints, 100 cases of commercial sexual exploitation of children, 112 cyber abuse complaints, and additional categories not directly relevant to this study. Over 2010–2025, the NCPA recorded more than 108,000 complaints in total. The UNICEF (2018) baseline study recorded internet use among 52.8% of young people, with average first access at age 13; among users aged 10–16, 46.3% reported communicating with strangers online and 27.9% reported meeting an online stranger in person. Save the Children research, reported by EconomyNext (2024), found that three in ten internet-using children in urban and semi-urban areas of Sri Lanka had been victimized by online perpetrators. The UN Women/UNFPA (2025) study recorded that 57% of respondents felt unsafe online and approximately one in three had personally experienced technology-facilitated gender-based violence; 64% of survivors subsequently stopped using social media, and 42% reported experiencing community rejection following online targeting. A separate clinical study recorded a 17.2% prevalence of internet addiction disorder among 15–19-year-olds in the Colombo District, with non-academic social media use identified as the leading associated factor (Ratnayake & Hewapathirana, 2023).

3.2 Online child sexual exploitation and abuse: global and domestic evidence

The WeProtect Global Alliance (2025) Global Threat Assessment, synthesising data from over 350 member organisations, recorded a 1,325% increase in reports of AI-generated child sexual abuse material to the U.S. National Center for Missing & Exploited Children (NCMEC) between 2023 and 2024, and approximately 100 reports of financial sexual extortion received by NCMEC per day during 2024. The assessment identifies encrypted messaging spaces, generative AI tools, and livestreaming as accelerating vectors, and states explicitly that global safeguards are failing to keep pace with the scale and sophistication of the threat. The Disrupting Harm research programme, conducted by ECPAT International, INTERPOL, and UNICEF Innocenti (2022) across thirteen countries in Eastern and Southern Africa and Southeast Asia, recorded that as many as 20% of children surveyed in some participating countries had been subjected to online child sexual exploitation or abuse in the preceding year, and that grooming interactions observed in the research could take as little as nineteen seconds to progress, with an average duration of 43–45 minutes from initial contact to a sexualised request.

Domestically, the NCPA operates an online CSAM reporting portal, launched jointly with the Internet Watch Foundation (IWF), Save the Children, and ChildFund Sri Lanka, through which reports are assessed by IWF analysts, forwarded to the UK National Crime Agency, and relayed via INTERPOL to law enforcement in the hosting country (EconomyNext, 2024). At the portal's launch, then-President Ranil Wickremasinghe stated that the country lacked the institutional capacity to resolve the problem within a year and that police staffing for women's and children's protection desks required strengthening. Separately, the NCPA's Anti-Trafficking Surveillance Unit has documented cases of online CSAM production by foreign nationals resident in Sri Lanka, referred to INTERPOL and foreign law enforcement, including a case in which a foreign national's conviction proceeded in the Negombo High Court for the sexual abuse of two boys combined with production of CSAM, and a second case in Anuradhapura resulting in deportation and prosecution abroad (ECPAT International, cited in ECPAT Global Monitoring Report: Sri Lanka).

3.3 Sri Lanka's legal and regulatory framework

Sri Lanka acceded to the Council of Europe Convention on Cybercrime (Budapest Convention) in 2005, the first South Asian state to do so, enabling passage of the Computer Crimes Act No. 24 of 2007. An independent legal-gap analysis commissioned by ECPAT Sri Lanka (Verité Research, 2019) found that Sri Lankan law does not codify online grooming as a distinct criminal offence; that the Penal Code's definition of a child is applied inconsistently across relevant provisions, creating gaps in coverage; that the exchange element central to commercial sexual exploitation is not clearly codified; and that the Penal Code does not address computer-generated or simulated CSAM, leaving AI-generated material — the specific category the WeProtect Global Alliance (2025) reports has increased 1,325% globally — outside clear domestic criminal coverage. The same analysis found that the age of consent is inferred, rather than explicitly stated, from statutory-rape provisions at 16 years, and that those provisions address only male offenders and female child victims, leaving other victim-offender configurations without equivalent statutory coverage. The Online Safety Act No. 9 of 2024 subsequently created an Online Safety Commission empowered to address categories including child abuse material, but the Commission answers to the President, remains unappointed as of the period covered by this study, and was not designed to resolve the minimum-age question presently before Parliament.

Table 1 summarises the quantitative evidence reviewed in Sections 3.1 and 3.2.

Source	Scope	Key finding	Year
NCPA	Sri Lanka, national	5,523 complaints Jan–Jun 2026 (100 commercial sexual exploitation, 112 cyber abuse); >108,000 complaints 2010–2025	2026
UNICEF	Sri Lanka, national	52.8% of youth online; avg. first access age 13; 46.3% of 10–16-y/o users contacted by strangers; 27.9% met a stranger in person	2018
Save the Children (in EconomyNext)	Urban/semi-urban Sri Lanka	3 in 10 internet-using children victimized by online perpetrators	2024
UN Women / UNFPA	Sri Lanka, national	57% felt unsafe online; ~1 in 3 experienced TF-GBV; 64% of survivors stopped using social media; 42% experienced community rejection	2025
Ratnayake & Hewapathirana	Colombo District, 15–19 y	17.2% prevalence of internet addiction disorder	2023
WeProtect Global Alliance	Global	1,325% rise in AI-generated CSAM reports to NCMEC (2023–2024); ~100 sextortion reports/day to NCMEC in 2024	2025
Disrupting Harm (ECPAT/INTERPOL/UNICEF Innocenti)	13 countries, E./S. Africa & SE Asia	Up to 20% of children subjected to OCSEA in past year; grooming as fast as 19 seconds, average 43–45 minutes	2022

Table 1. Summary of institutional, survey, and global threat-assessment evidence.

3.4 Original investigation and third-party platform-accountability findings

The author's OSINT investigation, submitted to Sri Lankan authorities — including the NCPA and the police — and to NCMEC in 2026, recorded the absence of functioning identity or age verification on the networks and accounts examined. Independently, Hattotuwa (2026a) recorded ten Facebook accounts using AI-generated identities of women, publishing 1,706 posts between February 2025 and June 2026 that drew 890,632 engagements; 113 clusters of identical captions and 676 near-duplicate post pairs across five coordinated accounts; conversion links routed through bios, shortlinks, and first comments to adult-chat, gambling, and cryptocurrency destinations; and foreign-language variants of the same architecture (Spanish, Italian, Malayalam) achieving more than three times the per-post engagement efficiency of the Sri Lankan set. This report concerns adults targeted for financial and romantic-scam purposes, not children, and is not evidence of child exploitation; its relevance is that it independently demonstrates, over a sixteen-month period and under close outside scrutiny, that Meta's transparency and enforcement tools fail to surface even large, coordinated, anonymity-dependent deception operating inside Sri Lanka's own user base.

3.5 Neurodevelopmental and trauma-psychology literature

Paulus et al. (2019), using ABCD Study neuroimaging data, recorded associations between screen media activity and measurable differences in cortical thickness, sulcal depth, and grey matter volume in developing brains. Hutton et al. (2020) recorded associations between higher screen time and reduced microstructural integrity of white matter in regions supporting language and executive function. Maza et al. (2023) recorded associations between habitual social-media checking behaviour and longitudinal changes in functional brain development within the same cohort. Separately, and foundationally for the trauma literature rather than the neuroimaging literature, Finkelhor and Browne (1985) identified four traumagenic dynamics produced by child sexual abuse — traumatic sexualization, betrayal, powerlessness, and stigmatization — and modelled these as distorting a child's cognitive self-concept and worldview. Finkelhor and Browne's own 1985 paper models the resulting injury; it does not itself propose that this injury cycles back into renewed exposure. That closing step is addressed by a distinct, dedicated literature: Classen, Palesh, and Aggarwal's (2005) review of approximately 90 empirical studies on sexual revictimization found that roughly two in three individuals sexually victimized in childhood are revictimized later, with the severity of the original abuse among the best-documented predictors of subsequent revictimization. Section 4.2 draws on Finkelhor and Browne (1985) for the injury this article's model calls Node 2, and on Classen et al. (2005) for the revictimization link the model calls Node 4 — the loop connecting the two, from injury through altered risk perception back to renewed vulnerability, is this article's own synthesis of both literatures, not a claim made by either source individually.

3.6 Domestic ban-mechanism policy record

Sundararaj et al. (2025), reviewing the neurodevelopmental and psychosocial evidence for a Sri Lankan clinical audience, recommended stricter content moderation, digital literacy programming, and parental controls, and did not recommend a blanket age ban, citing risk of displacement into unregulated spaces and the role of social media as a support channel for marginalised groups including LGBTQIA+ adolescents. Barnes et al. (2026), in an early observational evaluation of Australia's under-16 ban published in The BMJ, surveyed 408 adolescents aged 12–17 in New South Wales just before the law took effect and again three months later, finding more than 85% of under-16 respondents still using restricted platforms at follow-up, mostly through their own accounts (54–67%), with 66% reporting exposure to a platform age-verification check, most commonly a self-declared age or an uploaded photo. The study's own authors state the sample was small, drawn from a single state, and relied on self-report, which may limit accuracy and generalisability, though they report the finding held under sensitivity analysis. The accompanying BMJ editorial (Purba, 2026) and independent expert commentary (Science Media Centre,

2026) caution against reading the result as proof the policy has failed: both stress that three months is too short a window to separate weak implementation from a genuinely ineffective mechanism, and that wide confidence intervals around a sample this size cannot rule out modest real effects. This article treats the Barnes et al. (2026) finding — also reported in Sri Lanka by Hattotuwa (2026b) alongside Australian eSafety Commissioner data recording approximately 4.7 million accounts removed or restricted by mid-December 2025 and separate Commissioner reporting that roughly seven in ten parents of previously restricted-platform-using children say their child retained an account, as an early and contested signal of weak early-stage compliance rather than a settled verdict on the mechanism. Hattotuwa's policy brief to the Prime Minister's office (2026c) additionally reported LIRNEasia survey data recording internet access in 34% of Sri Lankan households with children nationally and 21% among the poorest households, with 23% of children living in a household with a computer and 28% of parents computer-literate; the national 34% figure is independently confirmed in LIRNEasia's own published survey brief (LIRNEasia, 2020).

Table 2 situates Sri Lanka's proposed mechanism against comparative international approaches, drawing on Sundararaj et al. (2025) and the sources in Section 3.6. The mechanisms listed are not points on a single spectrum toward Sri Lanka's proposal: parental consent, usage-time limits, and identity verification are structurally different policy instruments addressing different failure modes, and are listed together here only to map the range of approaches states have taken, not to imply they are interchangeable or that any one is a proven default. South Korea's entry requires a specific caveat, addressed fully in Section 4.5.2: the general real-name verification statute summarised here was struck down as unconstitutional by South Korea's Constitutional Court in 2012.

Jurisdiction	Mechanism	Age threshold	Verification method
Australia	Platform-mandated exclusion	Under 16	Self-declared birthdate / platform-side age estimation
United Kingdom	Design restrictions + proposed exclusion	Under 16 (proposed)	Not yet finalised
France	Parental consent requirement	Under 15	Parental consent
United States / European Union (COPPA-model)	Parental consent requirement	Under 13	Parental consent
China	Usage-time restriction	Under 18	Real-name registration + 60-minute daily limit

Jurisdiction	Mechanism	Age threshold	Verification method
South Korea	Mandatory real-name identity verification (general statute struck down 2012)	All users (as originally enacted)	Government-linked identity verification — ruled unconstitutional by the Constitutional Court of Korea, Aug. 2012
Sri Lanka (proposed)	Under review: exclusion vs. restriction	Under 16 (bill); under 12 (earlier proposal)	Undetermined; no verification framework yet enacted

Table 2. Comparative international regulatory approaches to minors' social media access, adapted from Sundararaj et al. (2025) and Sections 3.2–3.6. Entries are not mutually comparable policy instruments; see caveat above.

4. Discussion

4.1 Interpretation in relation to the thesis

The results in Sections 3.1 through 3.4 establish that online harm to Sri Lankan minors is extensive, active, ongoing, and in its most serious forms criminal, corroborated across independent institutional, global threat-assessment, legal, and investigative sources, including the author's own reported findings. This satisfies the first half of the thesis without qualification: the harm is real, urgent, and not dependent on any single contested data source. Section 3.3 further establishes that this harm occurs inside a documented legal gap — grooming uncodified, computer-generated CSAM unaddressed, an Online Safety Commission that remains unappointed — meaning that even where conduct is identified, prosecution is not assured under current statute. The results in Section 3.5 establish a plausible, though not fully settled, biological and psychological mechanism connecting exposure to structural and functional change in the developing brain and to compounding trauma, consistent with, rather than independent of, the harm identified in Sections 3.1–3.4. The results in Section 3.6, however, do not support the specific ban mechanism currently before Parliament, though the strength of that conclusion should be stated precisely. The Australian evaluation is the most directly comparable enforcement precedent available, and it records a weak early compliance signal — the large majority of under-16s retained platform access at three-month follow-up — in a jurisdiction with a dedicated independent regulator and greater enforcement resourcing than Sri Lanka currently possesses. The study's own authors and its accompanying editorial caution that a small, single-state, self-reported sample over a three-month window cannot distinguish weak implementation from a genuinely ineffective mechanism, so this article treats the finding as evidence that a self-declared-age ban is not a proven working mechanism, rather than as proof that identity-verified access limitation cannot work. That distinction matters for the argument that follows: it is what makes Track B a claim about a different, untested

mechanism rather than a claim that the Australian result has already been generalised to it. Combined with the LIRNEasia infrastructure data in Table 2, which indicates that Sri Lankan children access the internet predominantly through shared rather than personal devices, the self-declared-age mechanism the current bill relies on is mismatched to the environment it would need to operate in twice over.

4.2 The recursive mechanism of harm

The evidence in Sections 3.2 and 3.5 supports a specific structural claim, illustrated in Figure 1: the harm documented in this study behaves as a self-reinforcing loop rather than a series of independent incidents. The loop as a whole — the return arrow from Node 4 back to Node 1 — is this article's own synthesis of two separate literatures; no single cited source proposes the full four-node cycle, and the paragraph below is explicit about which claim comes from which source. Unverified exposure (Node 1) — online contact between a child and an unidentified adult, unmediated by any adult or institutional gatekeeper — is the entry point the Disrupting Harm data show can convert to a sexualised request within as little as nineteen seconds (ECPAT International, INTERPOL, & UNICEF Innocenti, 2022). Where this contact results in sexual exploitation or abuse, Finkelhor and Browne's (1985) traumagenic dynamics model (Node 2) provides the clinical account of the resulting injury: traumatic sexualization, betrayal, powerlessness, and stigmatization, operating together to distort a child's self-concept and worldview. Finkelhor and Browne's model addresses this injury and its effect on self-concept; it stops there. The next step — that this injury recalibrates a child's subsequent pattern-recognition and threat appraisal (Node 3) — draws on the broader trauma and neurodevelopmental literature reviewed in Section 3.5 rather than on Finkelhor and Browne specifically. The final step, closing the loop, is supported by dedicated revictimization research rather than by the trauma-dynamics or neurodevelopmental literature: Classen, Palesh, and Aggarwal's (2005) review of approximately 90 empirical studies found that roughly two in three children sexually victimized are revictimized later, establishing increased vulnerability to repeat victimization (Node 4) as an empirically supported, though probabilistic rather than universal, return to a state of renewed exposure.

This structure matters for policy because it identifies exactly two points at which external intervention can break the cycle rather than merely respond to one iteration of it. Track A intervenes at Node 1 from the perpetrator's side: identifying, prosecuting, and removing the adult responsible, which requires the strengthened statutory basis discussed in Section 3.3 and detailed in Section 4.5.1. Track B intervenes at Node 1 from the contact-mechanism side: removing the anonymity that allows an unverified adult to reach a child at all, detailed in Section 4.5.2. Both tracks target the same node because Node 1 is the only point in the loop external policy can act on directly; Nodes 2 through 4 are internal psychological and behavioural

sequelae that clinical intervention, not platform or criminal-justice policy, must address. A policy response confined to age exclusion — asking a user to self-report being under 16 — does not credibly interrupt Node 1, as the Australian compliance data in Section 3.6 demonstrate; it neither identifies the offender (Track A) nor establishes who is actually contacting the child (Track B).

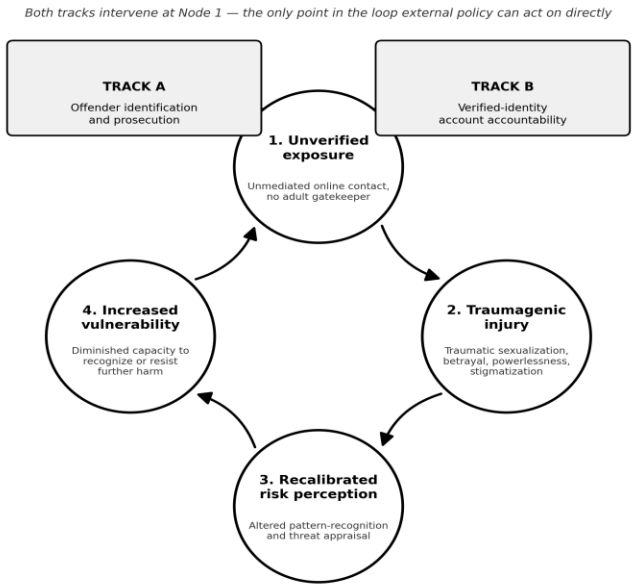

Figure 1. The recursive harm loop (author's synthesis), combining Finkelhor and Browne's (1985) traumagenic dynamics model (Node 2) and Classen, Palesh, and Aggarwal's (2005) revictimization findings (Node 4), with its two external intervention points.

4.3 Limitations

This study has seven principal limitations. First, the author's own OSINT investigation is reported here only at the level of its structural finding, with methods and operational detail withheld for the protective reasons stated in Section 2; this prevents independent replication of that component and means its inclusion rests on the author's submission to Sri Lankan authorities and NCMEC rather than on peer review. Second, the study is documentary and comparative rather than statistical: no new survey, experiment, or dataset was generated, and triangulation across independently collected sources cannot fully substitute for a single nationally representative Sri Lankan OCSEA prevalence study, which does not yet exist. Third, the neurodevelopmental literature reviewed in Section 3.5 is drawn predominantly from a United States cohort (ABCD Study); its applicability to Sri Lankan adolescents,

who differ in socioeconomic profile, device access, and cultural context, has not been independently established. Fourth, the Disrupting Harm data in Section 3.2, while methodologically the strongest cross-national OCSEA evidence available, covers thirteen countries in Eastern and Southern Africa and Southeast Asia and does not include Sri Lanka directly; its figures are used here as a regional and structural benchmark, not a Sri Lanka-specific prevalence estimate. Fifth, the comparative enforcement evidence in Section 3.6 concerns Australia and, secondarily, the European Union; both differ from Sri Lanka in regulatory capacity, per-capita income, and device-ownership patterns, and the direction of any bias this introduces cannot be determined from the present evidence. Sixth, whether Track B's platform- or third-party-held verification design would itself survive constitutional or rights-based challenge in Sri Lanka is not established by this article; Section 4.5.2 distinguishes the design from South Korea's invalidated statute on the reasoning available, but that distinction has not been tested against Sri Lankan constitutional law or evaluated by a legal scholar. Seventh, Track B has not been evaluated against children's own participation rights under UNCRC Article 12 and General Comment No. 25, as Section 4.5.2 notes; this article identifies the gap but does not close it.

4.4 Directions for future research

Four research priorities follow directly from these limitations. A nationally representative baseline study of Sri Lankan child internet use and OCSEA prevalence specifically, spanning all provinces and language communities and updating the 2018 UNICEF figures, is the single highest-priority gap and is independently recommended by Hattotuwa (2026c). A Sri Lanka-specific replication of the Disrupting Harm methodology would close the gap identified in Section 4.3 and allow direct regional comparison. Sri Lanka-specific neurodevelopmental research, even at a modest pilot scale, would establish whether the ABCD Study findings generalise to the local population or require local qualification. Finally, prospective evaluation of verified-identity account accountability against the status-quo anonymity model, using prosecution rates and repeat-victimization rates as outcome measures rather than platform-reported compliance alone, would generate the enforcement-outcome evidence that Section 3.6 shows is currently available only for the ban mechanism and not for its proposed alternative.

4.5 Two Tracks: Offender Accountability and Enforceable Access Limitation

The loop modelled in Section 4.2 has exactly one external entry point, Node 1, and the evidence in Sections 3.2 through 3.4 indicates that two concurrent interventions are required to close it, addressing the perpetrator side and the contact-mechanism side respectively. Neither substitutes for the other, and neither is satisfied by a blanket age-declaration ban.

4.5.1 Track A: Offender identification and prosecution

The individuals responsible for grooming, sextortion, and the production or circulation of CSAM documented in Section 3.2 are committing criminal acts under international instruments Sri Lanka has ratified — the Budapest Convention since 2005, and the Optional Protocol on the Sale of Children, Child Prostitution and Child Pornography — and this obligation is independent of any platform-access policy. Section 3.3 shows that domestic statute has not yet caught up with this obligation: grooming is not a codified offence, computer-generated CSAM falls outside clear Penal Code coverage, and the Online Safety Commission empowered to address these categories remains unappointed. Closing this gap requires four concrete steps, each with a distinct owner. First, the Ministry of Justice, drawing on the specific provisions recommended by Verité Research (2019), should draft Penal Code and Computer Crimes Act amendments codifying grooming as a standalone offence, extending statutory coverage to computer-generated and AI-mediated CSAM, and correcting the victim-offender coverage gap Section 3.3 identifies. This article does not have a sourced basis for how long such drafting would take in Sri Lanka specifically and does not assert one; what can be said is that drafting can begin immediately, since the gap analysis already exists, and that the timeline is a matter for the Ministry of Justice to specify rather than for this article to estimate. Second, Parliament should appoint the Online Safety Commission provided for under the 2024 Act and fund it at a level proportional to the complaint volume in Table 1 — 5,523 complaints in six months implies a caseload the NCPA's current resourcing was not built for, and the Commission's enabling statute already exists, so this step requires appointment and budget allocation rather than new legislation. Third, the NCPA's Anti-Trafficking Surveillance Unit and the existing NCPA–IWF–INTERPOL reporting pipeline documented in Section 3.2 should receive dedicated cyber-forensic staffing, addressing the capacity gap then-President Wickremasinghe acknowledged at the reporting portal's 2024 launch. Fourth, prosecutorial guidance and judicial training on the amended provisions should follow enactment, so that a codified offence translates into charges rather than sitting alongside the uncharged conduct Section 3.3 describes. None of this depends on, or is satisfied by, a minimum-age law; a perpetrator identified and prosecuted under a strengthened statute is removed from Node 1 regardless of what age threshold a platform enforces.

4.5.2 Track B: Verified-identity account accountability

What made the harm documented in Section 3.2 and 3.4 possible was not the absence of an age gate; in both the author's investigation and Hattotuwa's (2026a) independent findings, it was the absence of any verified identity behind the contacting account. A predator does not need to defeat an age check if the account approaching a child answers to nobody. This suggests a mechanism distinct from the ban currently before Parliament: mandatory identity verification, applied to every account on a platform as a condition of access, rather than an age threshold applied only to minors.

A bank account, a driving licence, and a mobile SIM registration already require verified identity in Sri Lanka precisely because the harms possible through those instruments warrant it; a social media account through which an adult can contact a child directly is not self-evidently a lower-stakes instrument.

The most directly relevant international precedent for this proposal is not a success story and must be addressed on that basis. South Korea operated a mandatory real-name verification law for major websites from 2007, requiring sites with more than 100,000 average daily visitors to record posters' identities, typically via national resident registration numbers. In August 2012 the Constitutional Court of Korea struck the law down unanimously, holding that it violated users' freedom of expression and their right to informational self-determination. The Court found that the law had not produced a measurable decline in the abusive postings it targeted, had pushed Korean users toward offshore platforms outside domestic jurisdiction, and had concentrated identity data in centralised databases that were subsequently breached, including a 2011 breach exposing 35 million records — more than half the national population. This is a direct precedent for exactly the harm Track B must avoid, not a demonstration that identity verification works.

Track B's design differs from the invalidated Korean statute in three respects, each tied to a specific element of the Court's reasoning. First, it does not concentrate identity in a state-accessible database: the Korean law tied verification to a state-issued registration number retained by platforms, while Track B specifies platform- or third-party-held credentials never transmitted to government. Second, it is not general-purpose content-based speech regulation: the Korean law's stated purpose was curbing defamatory and abusive postings generally, which the Court found insufficiently served the restriction on expression it imposed, while Track B's purpose is restricting unverified adult-to-child contact specifically, a narrower harm with the direct evidentiary basis set out in Section 3.2. Third, it does not eliminate anonymous expression as such, since accountability attaches to contact-enabling account activity rather than to posting or speech content generally. Independent technical analysis of privacy-preserving verification design supports the direction of this distinction without resolving it: a 2025 review by the privacy-focused browser company Brave of zero-knowledge-proof and similar credential-based age- and identity-verification schemes found that centralisation and exclusion risk depends heavily on implementation choices — specifically, how many entities can issue or hold the underlying credential, and whether verification events are linkable across sessions — rather than being an unavoidable property of verification itself (Brave, 2025). That review is industry technical commentary, not constitutional or Sri Lankan legal analysis, and does not itself establish that Track B's specific design would satisfy a Sri Lankan court applying reasoning similar to the Korean Constitutional Court's. Whether this distinction would survive constitutional or rights-based challenge in Sri

Lanka is not established by this article and is flagged as a research and legal-drafting priority in Section 4.4 and as Limitation 6 in Section 4.3, not asserted as settled.

A second objection, raised in Section 3.6 by Hattotuwa (2026b) but not yet addressed in this article, concerns children's own participation rights rather than institutional surveillance risk. Article 12 of the UN Convention on the Rights of the Child, and General Comment No. 25 extending it to the digital environment, require that children's own views be heard in decisions affecting them online, including decisions to restrict their access. Track B is an access-control mechanism and has not, in the drafting reflected in this article, been evaluated against that standard. A revision-stage consultation with adolescent stakeholders would close this gap; Amnesty International's June 2026 gathering of Sri Lankan young people's perspectives, reported in the Daily Mirror's coverage of the ban debate, is cited here only as an illustration that such consultations are already happening in this policy debate and are a methodology this article's authors did not themselves undertake, not as a specific protocol this article endorses. That gap is listed as Limitation 7 in Section 4.3.

This recommendation must also be reconciled with the institutional-surveillance objection discussed in 4.1: that an identity-verification regime, in a state whose Online Safety Commission remains unappointed and whose recent legislative history includes the contested Online Safety Act, risks becoming an instrument of political surveillance rather than child protection. That risk is real and is not dismissed here; it is an argument about how identity verification is implemented, not an argument against verification itself. The mechanism Hattotuwa (2026b) objects to is a state-linked age-gate checking an NIC, driving licence, or passport against every user to determine who is under 16, creating a de facto national identity register — structurally close to the Korean statute the Constitutional Court invalidated. The mechanism proposed here differs structurally: platform- or third-party-held identity, tied to account accountability for all users rather than age exclusion for a subset of them, with the verifying credential never transmitted to or stored by any government body, no persistent link between a verified identity and a state database, and disclosure to law enforcement restricted to judicial order in the manner already governing banking and telecom records. Under this design, the platform gains the capability, already normal in banking and telecommunications, to say who a given account belongs to when that account causes harm; the state does not gain a new surveillance capability.

This design also answers the enforcement-failure evidence in Section 3.6 more directly than an age ban does. The Australian evaluation found weak early compliance in part because a self-declared birthdate is trivial to falsify while the underlying account remained anonymous regardless of the date entered; verifying who a person is addresses that specific failure mode in a way that verifying how old they claim to

be does not, though this article treats that as a reasoned expectation rather than a demonstrated result, since no jurisdiction has yet evaluated this precise mechanism against this precise outcome. It does not depend on the shared-device problem raised by the LIRNEasia data, since identity verification attaches to the account rather than the device, and a shared family phone can hold multiple verified accounts exactly as it can hold multiple verified banking apps today. The recommendation carries its own precondition, consistent with the limitation stated in 4.3: it requires the independent oversight body and data-protection safeguards Sri Lanka does not yet have, and building the mechanism ahead of that institutional capacity would recreate the surveillance risk this section sets out to avoid. Sequencing — institutional capacity first, verification infrastructure second — is as important to this recommendation as the mechanism itself.

4.6 Synthesis

Sri Lanka is not a special case; it is a visible one. A culturally conservative, family-centred society confronting an unmediated global contact channel present in every child's pocket makes the contradiction harder to hide here than in countries better practised at not examining it. The mechanism producing this harm — a developing brain and a developing sense of self, a contact channel indifferent to the identity of either party, and the near-total removal, within roughly fifteen years, of mediating structures human societies relied on for millennia — is not a Sri Lankan phenomenon. It is the condition every country with meaningful smartphone penetration now lives inside. Sri Lanka's own institutions and researchers have already done much of the work this article draws on: the NCPA has itself flagged its resourcing gap to Parliament, Verité Research mapped the legislative gap for ECPAT Sri Lanka years before this article was written, and domestic researchers have independently stress-tested the ban mechanism against the same international evidence reviewed here. This article's contribution is to bring that domestic work together with the global evidence base and show that it points toward one answer. The evidence assembled here supports urgent, two-track action and does not support a blanket age ban as a sufficient instrument of either track: prosecuting identified offenders under a strengthened legal framework, and mandatory verified-identity accountability applied to every user, sequenced behind the independent oversight body Sri Lanka does not yet have. A ban that excludes on the basis of an unverifiable self-declared birthdate protects no one it claims to protect, in Sri Lanka or anywhere else it has been tried. Both gaps identified here — the prosecutorial one and the identity one — are closable largely with tools and statutory language Sri Lanka already has in draft or in hand; what remains is sequencing, appointment, and funding rather than a search for new solutions. This article specifies exactly where to close them.

Funding

This research received no external funding.

CRediT Author Statement

Anthony Vondoom: Conceptualization, Investigation, Writing – original draft, review & editing.

AI Disclosure Statement

Generative AI tools were used under the author's direction to support literature search, source triangulation, citation formatting, and structural editing of drafts written by the author. All factual claims, statistics, and citations were reviewed and checked by the author against primary or authoritative secondary sources before inclusion. Two structured AI-assisted review passes were conducted specifically to stress-test the manuscript's evidentiary claims and citation accuracy prior to finalisation; issues identified through this process, including corrections to how comparative international precedents were characterised, were independently verified and incorporated by the author. The author takes full responsibility for the content, accuracy, and conclusions of this article.

Acknowledgements

The author is grateful to the Sri Lankan child protection professionals, investigators, and civil society researchers who are, at the time of writing, actively working these cases on the ground — often without adequate resourcing or public recognition. Their work identifying, reporting, and pursuing accountability for online child exploitation informed and motivated this article, THANK YOU ALL for standing up for what is right.

References

Barnes, C., Hall, A., Mantach, S., Oldmeadow, C., Attia, J., Backholer, K., Williams, C., Tefera, Y., Kay, F., & Wolfenden, L. (2026). Assessing early effects of Australia's Social Media Minimum Age Act on adolescents' social media use: Observational study. *The BMJ*, 393, bmj-2026-363695.

Brave. (2025, November 5). The limits of zero-knowledge proofs for age verification. Brave Blog.

Classen, C. C., Palesh, O. G., & Aggarwal, R. (2005). Sexual revictimization: A review of the empirical literature. *Trauma, Violence, & Abuse*, 6(2), 103–129.

Constitutional Court of Korea. (2012, August 23). Online Real-Name Case, 2010Hun-Ma47 [invalidating Article 44-5 of the Act on Promotion of Information and Communications Network Utilization and Information Protection].

ECPAT International, INTERPOL, & UNICEF Office of Research – Innocenti. (2022). Disrupting Harm: Evidence on online child sexual exploitation and abuse across 13 countries in Eastern and Southern Africa and Southeast Asia.

ECPAT International. (n.d.). Global Monitoring Report on the status of action against commercial sexual exploitation of children: Sri Lanka.

EconomyNext. (2024, March 29). Sri Lanka launches online child sexual abuse reporting portal.

Finkelhor, D., & Browne, A. (1985). The traumatic impact of child sexual abuse: A conceptualization. *American Journal of Orthopsychiatry*, 55(4), 530–541.

Hattotuwa, S. (2026a, July 7). AI profiles of women on Facebook: Identity deception, sexualised engagement, and platform accountability. Centre for International Reconciliation / sanjanah.wordpress.com.

Hattotuwa, S. (2026b, July 14). Social media bans: Protecting children or building a panopticon. *Newswire*; reposted sanjanah.wordpress.com.

Hattotuwa, S. (2026c, July 30). Children, harms, and rights: Policy brief for Prime Minister on proposed social media restrictions for under-16s in Sri Lanka. sanjanah.wordpress.com.

Hutton, J. S., Dudley, J., Horowitz-Kraus, T., DeWitt, T., & Holland, S. K. (2020). Associations between screen-based media use and brain white matter integrity in preschool-aged children. *JAMA Pediatrics*, 174(1), e193869.

LIRNEasia. (2020, May 20). Internet access for households with children under 18 in Sri Lanka is 34%. LIRNEasia AfterAccess survey brief.

Maza, M. T., Fox, K. A., Kwon, S., et al. (2023). Association of habitual checking behaviours on social media with longitudinal functional brain development. *JAMA Pediatrics*, 177(2), 160–167.

National Child Protection Authority, Sri Lanka. Child abuse complaint statistics, January–June 2026, and 2010–2025 decadal figures, as reported to Parliament.

Paulus, M. P., Squeglia, L. M., Bagot, K., Jacobus, J., Kuplicki, R., Breslin, F. J., et al. (2019). Screen media activity and brain structure in youth: Evidence for diverse structural correlation networks from the ABCD study. *NeuroImage*, 185, 140–153.

Purba, A. K. (2026). Learning from Australia's social media age restriction policy [Editorial]. *The BMJ*, 393, bmj-2026-100017.

Ratnayake, C., & Hewapathirana, C. (2023). Internet addiction disorder and its associated factors among 15–19-year adolescents in Colombo District, Sri Lanka. *Ceylon Journal of Science*, 52(4), 1–10.

Science Media Centre. (2026, June 24). Expert reaction to study assessing early evidence of association between Australia's Social Media Minimum Age Act and adolescents' social media use.

Sundararaj, R., Udakumbura, N., Rathnayake, L., & Chandradasa, M. (2025). Scroll control: Should Sri Lanka enforce a social media ban for minors? *Sri Lanka Journal of Psychiatry*, 16(2), 50–52.

UN Women Asia-Pacific & UNFPA. (2025, September). Research study on technology-facilitated violence experienced by women and marginalized groups in Sri Lanka.

UNICEF Sri Lanka. (2018). Keeping children safe and empowered online: A study on internet use and online safety of children in Sri Lanka.

Verité Research. (2019). Sri Lanka: Online child sexual exploitation legal gap analysis. Commissioned by ECPAT Sri Lanka (PEaCE).

WeProtect Global Alliance. (2025, December 11). Global Threat Assessment 2025. Produced with the CPC Learning Network, Columbia University Mailman School of Public Health.